IMPACT OF LOGISTICS ON READINESS AND LIFE CYCLE COST: A DESIGN OF EXPERIMENTS APPROACH

Keebom Kang

Naval Postgraduate School
Graduate School of Business and Public Policy
Monterey, CA 93933, USA

Mary McDonald

Naval Postgraduate School
Operations Research Department
Monterey, CA 93933, USA

ABSTRACT

In this paper we develop two models that can be used to identify critical logistics factors that impact military readiness and the life cycle cost. The first one, a discrete-event simulation model, estimates the operational availability of a weapon system given input parameters under a certain scenario. The second one, a spreadsheet model, computes the life cycle cost using the same input parameters for the simulation model. Our approach is intended to serve as a basis for discussion between program offices concerned with cost and operational commands concerned with operational availability.

1 INTRODUCTION

The development of modeling and simulation (M&S) to support Life Cycle Management is complex and requires a close examination of the factors involved and results obtained. The insights derived from M&S can contribute significantly to operational readiness and the cost of readiness. Consequently, research is needed to identify important data and key factors that affect Life Cycle Management.

In this paper we develop two models that can be used to identify critical factors that impact military readiness and the life cycle cost. The first one is a simulation model using the Arena[®] simulation language (Kelton et al. 2010). This model estimates the operational availability of a major weapon system given input parameters under a certain scenario. The second one is an Excel[®] spreadsheet based life cycle cost model that computes the life cycle cost using the same input parameters for the simulation model. Although the life cycle cost model can be embedded into the simulation model, we intentionally separated it utilizing the Excel spreadsheet, so that it can be used by financial managers to conduct what-if analyses virtually on any personal computer without obtaining a license or learning the simulation package.

We have selected the Light Armored Vehicle with a 25 mm Gun System (LAV-25) for our analysis. Our scenario includes 76 LAV-25s to be deployed with the Marine Expeditionary Forces. We are interested in readiness analysis for the LAV-25 fleet and in estimating the life cycle cost over a period of time (typically 20 years). In this paper we only consider operations and maintenance cost without including infrastructure, personnel or overhead costs. The models discussed in this paper can be applicable to any major weapon systems by providing new data sets. [See studies on unmanned aerial vehicles in Kang, Doerr and Sanchez (2006) and Kang et al. (2010).]

2 BACKGROUND

The goal of military logistics support is to maintain the highest possible level of readiness, commonly expressed as operational availability:

$$Ao = \frac{MTBM}{MTBM + MDT} = \frac{uptime}{uptime + downtime}$$

where MTBM is the mean time between maintenance, and MDT is the maintenance down time, which includes repair time and administrative and logistics delay times. Intuitively, operational availability is the fraction of time a weapon system is operational or mission capable. Clearly, operational availability can be improved by increasing MTBM (i.e., increasing reliability) and/or decreasing MDT (i.e., reducing repair or cycle-time). Thus, the two key issues to improve weapon systems readiness are reliability improvement and cycle-time reduction.

Reliability improvement and cycle-time reduction not only improve readiness but also save money. Better reliability requires less maintenance and fewer spare parts. The relationship between repair or cycle-time and inventory levels is critically important (yet, troublesome) in the military because it crosses physical, organizational, and financial barriers. Financial managers strive to consolidate and minimize stocks of piece-parts to free-up resources for other priorities. They also seek to get quick turnaround on repairable components in order to minimize pipeline inventory. However, stockout of spare parts or consumable components results in delays in repair processes and eventually serious readiness degradation. Cycle-time reduction in a military logistics channel (repair depots, intermediate-level maintenance, inventory control points, and supply centers) also means that more weapon systems are available in the field or fleet. On the other hand, increased cycle-time causes a vicious cycle of deteriorating military readiness. For instance, poor logistics support (e.g., lack of spare parts, personnel, and/or training) increases the cycle-time, which in turn decreases readiness, *Ao*. Increasing operational tempo (Op Tempo) also negatively affects readiness by generating more system failures, and increases the cost to support longer operating hours.

Understanding both military readiness and its implied cost is important to both warfighters and financial managers. There are many factors in the military logistics system that affect weapon system readiness and its life cycle cost. Clearly, simulation models of even relatively simple logistics systems can have a very large number of inputs—many of which may be uncertain or unknown—that potentially impact the model's performance. Some of these inputs are reliability, maintainability, logistics delay, repair turnaround time, sparing and operational tempo. In the design of experiments (DOE) literature, these are referred to as *factors*. In real-world experiments, it is difficult to control more than a handful of factors at a time. This is not the case for simulation experiments, where the analyst has the ability to specify the *levels* (values) for all of the input factors before running the simulation. Still, once the factors and potential levels have been determined, this creates a huge number of potential *scenarios* (or *design points*). For example, if an analyst wished to explore nine factors, each at 10 levels, there are one billion (10^9) different scenarios that could be considered. The design might need to be replicated for stochastic simulations, because specifying all input factors does not remove randomness from the output. Such a large experiment is clearly impractical. Even if it were possible to run all scenarios in a reasonable amount of time, the volumes of output data would easily overwhelm most post-processing analytic tools, leaving the analyst limited in their abilities to statistically interpret the results. Fortunately, efficient experimental designs can be used to specify a small number of suitable scenarios.

Kleijnen et al. (2005) discuss situations where various classes of designs are appropriate, but there is no one-fits-all design. In our explorations of readiness and life cycle cost analysis we want to screen many variables for importance, while simultaneously maintaining the ability to fit complex meta-models to a handful of input variables that are found to have the most impact on the responses. Given this, and the above design goals, the Nearly Orthogonal Latin Hypercube (NOLH) design constructed by Cioppa and Lucas (2007) is particularly useful.

3   MODELS

3.1   Simulation Model for Operational Availability

There are 76 LAV-25s to be deployed with the Marine Expeditionary Forces. We are interested in readiness analysis for the LAV-25 fleet and in estimating a life cycle cost over a period of time (typically 20 years). In this paper we only consider operations and maintenance cost without including infrastructure, personnel or overhead costs.

An LAV consists of 1,570 different parts. The data we collect from the Marine Corps Decision Support System (MCDSS) indicate that the Marine Corps has spent a total of $29 million to purchase all the parts to maintain LAVs in 2007 – 2009. After sorting the data in the order of the Extended Price (the total amount of money spent on each part; it is calculated by multiplying the part count by the unit cost), we note that more than $18 million dollars or approximately two-thirds of the total cost were spent on the five parts listed in Table 1. These parts are expensive and fail more often than others. Also these parts are critical, i.e., if any one of these fails, the LAV cannot operate.   We focus on these five parts for our analysis.

Table 1: The Five Major Components for LAV (in terms of the Extended Price)[Source: MCDSS 4.3.1.1, PartUsage_EO947 (2007-2009)]

	NSN	Part Name	Part Count	Unit Price	Extended Price *	Cumulative Extended Price
1	5860014304339	SENSOR UNIT,LASER	120	$89,794	$10,775,318	$10,775,318
2	1005015023970	CONTROL DISPLAY UNI	103	$27,683	$3,072,854	$13,848,172
3	2520219063912	DIFFERENTIAL,DRIVING	101	$22,475	$2,269,978	$16,118,151
4	2815014427645	ENGINE,DIESEL I	31	$41,757	$1,544,995	$17,663,146
5	2815015502615	ENGINE,DIESEL II	36	$26,890	$1,075,609	$18,738,755

*Extended Price = Part Count * Unit Price

In Table 2, the failure rate, the SL quantity, the criticality code and the level of repair are shown for each part. In our analysis, though, we did not use the failure rates in Table 2. Instead we allowed this parameter to vary over a fairly wide range in the design of experiments, using the NOLH. More details on the design of experiments will be presented in the following sections. The criticality code (Crit Code in Column 5) 5 indicates all these parts are critical; any one of these parts fails the LAV cannot operate. The SL quantity (Column 4) identifies the total number of part(s) that are required on a selected weapon system. The SL quantity of the driving differential (Part #3) is 4, which is interpreted as 4 serially connected driving differentials in an LAV. We are assuming that if any one of them fails, the LAV will not be operational. If the reliability of an individual driving differential is 0.90, the subsystem that consists of these 4 serially connected driving differentials is only 0.656 (or 0.90^4). Thus among these 5 parts, the driving differential (Part #3) is potentially more critical to operational availability of the LAV than other parts. The SL quantity of the laser sensor unit (Part 1) is 2, which means that two of these parts are serially connected.   The far right-hand-side column of Table 2 shows the level of repair for each component. Control display units (Part #2) and driving differentials (Part #3) are repaired at the intermediate-level (I-level) maintenance facility while laser sensor units (Part #1), and diesel engines (Parts #4 and #5) are maintained at the depot level (D-level). D-level as compared to I-level generally involves much longer turnaround times.

We consider the five critical components as shown in Table 1 for this study. When any of these parts fails, the faulty part is removed from the LAV, a spare is installed, and the faulty part is sent to the repair facility (I-level or D-level). After the repair is complete, the repaired part becomes a spare and is sent to

the spare pool. When a critical part fails, and a spare is not available, the LAV will be non-operational, or not mission capable, until a spare becomes available.

Table 2: Failure Rate, SL Quantity, Criticality Code and Level of Repair for Each Part [Source: Same as TABLE 1]

	Part Name	Failure Rate	SL Qty	Crit Code	Level of Repair
1	SENSOR UNIT,LASER	0.000211638	2	5	D-level
2	CONTROL DISPLAY UNIT	0.000363312	1	5	I-level
3	DIFFERENTIAL,DRIVING	8.90643E-05	4	5	I-level
4	ENGINE,DIESEL I	0.000109346	1	5	D-level
5	ENGINE,DIESEL II	0.000126983	1	5	D-level

The ranges of the failure rates (λ_i) of the individual parts are provided in Table 3, along with the ranges of the number of spare parts, repair turnaround times (I-TAT and D-TAT), and the operational tempo. The failure rate (λ_i) is expressed in terms of the number of failures per operating hour. The minimum λ_i value of 0.0001 failures per operating hour is equivalent to an MTBF of 25 years (= 1 / 0.0001 / 400) assuming an operational tempo of 400 hours per year. The maximum value of 0.005 failures per operating hour is equivalent to an MTBF of 0.5 years (= 1 / 0.005 / 400) at the same operational tempo. The spare levels are set at 1 through 10. The I-level and D-level turnaround times are set at 5 – 15 days and 30 – 60 days, respectively. Since the average operational tempo during the peace time is approximately 350 hours per year, and during the war time, is approximately 650 hours per year, we set the range of the operational tempo from 300 to 700 hours per year. We assume that the time between failures for each component follows an exponential distribution. Our simulation model estimates the average operational availability. Our goal is to better understand how changes in reliabilities (in terms of the failure rate, λ_i), number of spares for each part, repair turnaround times, and operational tempo impact operational availability and the life cycle cost.

Table 3. Ranges of Input Parameters

Input Parameter	Range
lambda i (λ_i, i = 1, 2, …, 5)	0.0001 – 0.005 failures per operating hr
spare i (i = 1, 2, …, 5)	1 - 10
I-TAT (I-level turnaround time)	5 – 15 days
D-TAT (D-level turnaround time)	30 – 60 days
Op Tempo	300 – 700 hrs/yr

Several designs are possible, but we use an NOLH with 257 design points (Cioppa and Lucas 2007) for our 13 factors. It can be easily constructed by entering the low and high values in Table 3 into a spreadsheet (Sanchez 2006) that is available from NPS's SEED Center for Data Farming web site (NPS 2010). For each scenario, the simulation model reads a row of data from the NOLH spreadsheet. The failure rates (λ_i, i = 1, 2, …,5) of the five parts are first read, followed by the number of spares for each part, the I-level turnaround time (I_TAT), the D-level turnaround time (D_TAT) and the operational tempo (Op_Temp). The λ_i values are expressed in terms of the number of failures per operating hour. The time between part failures is assumed to follow an exponential distribution for all 5 components we study. The repair turnaround times are assumed to follow symmetric triangular distributions with lower and upper bounds of 0.5(mean) and 1.5(mean), respectively.

The simulation model was developed using the Arena simulation language. A brief description of the model logic is as follows:

1. Read input data for each scenario.
2. Generate 76 LAVs.
3. Generate part failures. (Five failure times are generated. Whichever is the smallest value is the next failure time of the LAV.)
4. The faulty part is removed from the LAV
5. Update FMC count (fully mission capable LAVs): FMC = FMC – 1
6. If a spare part is available, install a spare and reset the FMC count: FMC = FMC +1. Otherwise wait in the queue until a spare is available from the spare pool until after Step 8.
7. The faulty part is sent to the repair facility.
8. After the delay as specified in the input, the repaired part joins the spare pool.
9. Repeat Steps #1 to #8 until the end of the simulation time.

At the end each of scenario, the simulation automatically calculates the operational availability: Ao = average FMC / total number of LAV. FMC is a time-persistent variable (see, e.g., Kelton et al. 2010) and the average value of FMC must be "time-averaged." Arena automatically computes the value.

We run a total of 257 scenarios, each of which is simulated over a period of 1,000,000 hours and the first 100,000 hours of observations were eliminated to remove initial bias. 900,000 hours of simulation is equivalent to approximately 7 replications of 20-year simulation. The average Ao (operational availability) from each scenario is automatically appended (see Figure 1, Column N) to the same input EXCEL worksheet that contains the factor values (simulation inputs). Columns A through M contain the inputs.

Figure 1. A Screen Shot of NOLH Input Parameters (Columns A through M) and Output Results (Ao in Column N and the life cycle cost in Column O)

3.2 Spreadsheet Model for Life Cycle Cost

The second model is the life cycle cost spreadsheet model that reads the same input scenarios and computes the life cycle cost over the next 20 years. The life cycle in this model only includes operations and maintenance costs (spare, repair, transportation, and operations costs).

Our scenario includes 76 LAVs. The operational tempo (Op Tempo or operating hours) changes as part of the design of experiment. We assume that the hourly operating cost to be $50, the hourly repair

cost, $300, and the transportation cost, $200 per failure. The annual capital discount rate of 7%, and the annual inventory rate of 20% are used. The life cycle is assumed to be 20 years for the LAV. Any of these values can be modified by the user.

We have developed a visual basic macro program to automate the life cycle cost computation for each of 257 scenarios. Once the macro is executed, it reads the input parameters in each scenario (row by row) in Figure 2 then writes them on the highlighted cells on the spreadsheet life cycle cost model. Once the spreadsheet model updates the life cycle cost, the result is written onto the same input worksheet (Column O in Figure 1). This procedure is repeated 257 times (once for each design point). When the macro execution is complete, the life cycle cost results are on Column O, and the operational availability results from the simulation are in Column N, along with input parameters in Columns A through M (see Figure 1). Then this worksheet is imported into the software JMP® (SAS 2008) for further analysis.

4   RESULTS

We begin assessing the output by examining the histograms of the simulation responses. This can be a way of "accidentally" performing verification and validation of a simulation model. Our results indicate that the average operational availability differs widely across the different scenarios, ranging from 0.573 to 0.995. The average Ao across the 257 scenarios is 0.830 with a standard deviation of 0.090. It appears that at least one of the input factors does, indeed, have a substantial influence on the system's performance (See Figure 2).

The average life cycle cost across the 257 scenarios is $52.1 million with a standard deviation of $14.6 million. The life cycle cost also differs widely across the different scenarios, ranging from $19.9 million to $99.9 million. It also indicates that at least one of the input factors does have a substantial influence on the life cycle cost (See Figure 3). The darkened areas of the Ao and cost distributions correspond to the subset of scenarios with lowest cost and highest Ao points, which will be explained with Figure 8.

Quantiles

100.0%	maximum	0.99520
99.5%		0.99348
97.5%		0.98177
90.0%		0.94315
75.0%	quartile	0.89972
50.0%	median	0.83795
25.0%	quartile	0.76866
10.0%		0.70506
2.5%		0.64539
0.5%		0.57274
0.0%	minimum	0.57256

Moments

Mean	0.8297314
Std Dev	0.0897837
Std Err Mean	0.0056006
Upper 95% Mean	0.8407604
Lower 95% Mean	0.8187023
N	257
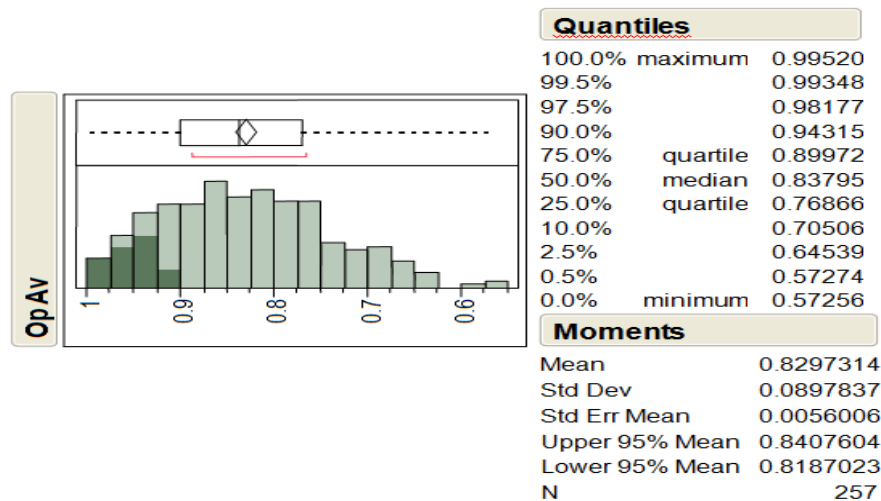
Figure 2. Distribution of Operational Availability (Op Av)

After confirming that the results appear reasonable, we turn to our main goals—identifying those factors and components that have the greatest impact on performance. A useful non-parametric data analysis technique is the regression tree, as shown in Figure 4. Regression trees are more human-readable and can be easier to describe than multiple regression models because they reveal the structure in the data in a simple way. Initially, the data are grouped in a single cluster. All potential input factors are examined to identify how best to split them to yield two leaves such that the variability in the response within each leaf decreases and the variability in the response between the leaves increases as much as possible.
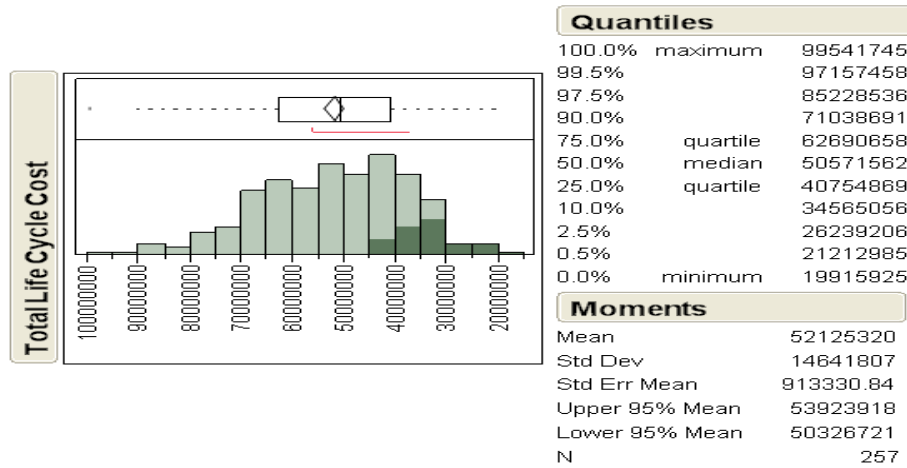
Figure 3. Distribution of Life Cycle Cost
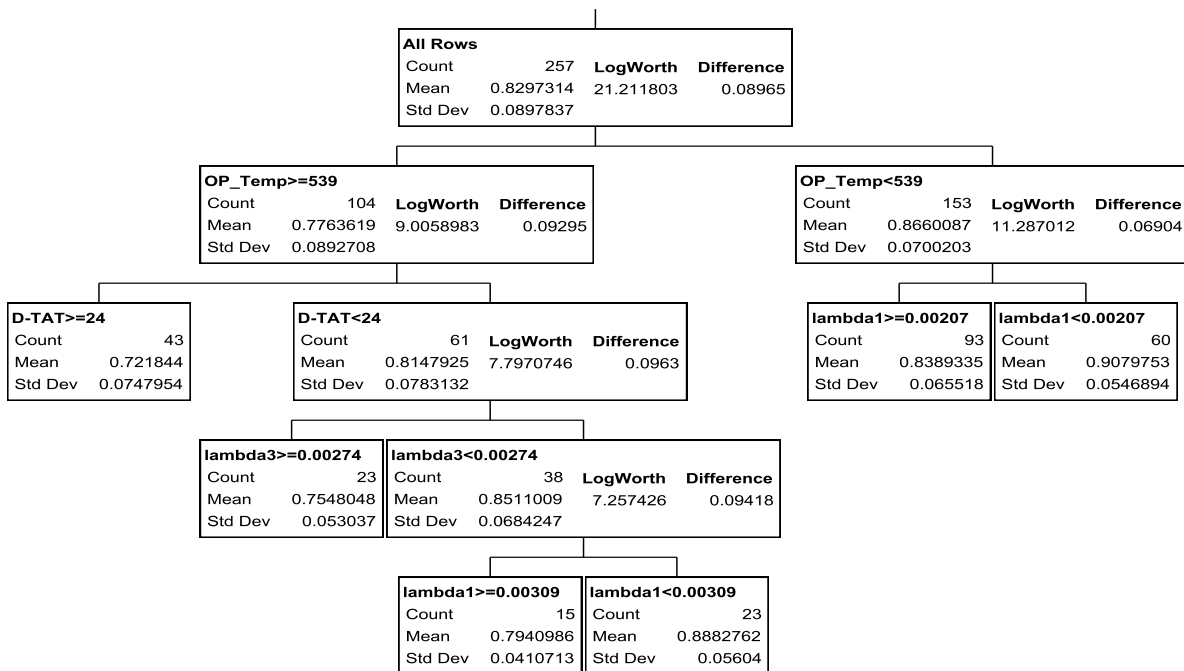
Figure 4. Regression Tree for the Average Ao

Figure 5 shows the regression tree for predicting the average Ao from the 257 simulation scenarios. The dominant factor is clearly the operational tempo (Op_Temp). For example, the first split at the top indicates that the average Ao is 0.776 across the 104 scenarios that have an operational tempo of 539 hours or more per year. In contrast, the average Ao is 0.866 (11.6% higher) among the 104 scenarios that had an operational tempo less than 539 hours per year. As the operational tempo increases (i.e., more operating hours), the more failures occur and obviously the operational availability goes down. In the second split when the operational tempo is greater than 539, the depot turnaround time (D-TAT) becomes a critical factor. Then the failure rates (lambda 1 and lambda 3) of the laser sensor unit (Parts #1) and the driving differentials (Part #3) become critical.

Regression trees are non-parametric approaches for fitting a statistical model to the simulation output. They can be good at identifying subsets of the output that behave much differently than the rest. Regression metamodels can also be valuable. They may confirm the regression tree results concerning which factor or factors have the greatest influence on the results, or they may allow more succinct descriptions of the simulation model's performance if it can be well-described by simple polynomial metamodels.

Accordingly, we fit regression metamodels of the Ao as a function of main effects and two-way interactions of the 13 input factors. After noticing that the interactions are not significant compared to the main effects, we develop a simple linear model with 13 main factors without interaction effects. The parameter estimates are shown in Figure 6. This model had an R^2 of 0.90. Large |t_ratio|s for the OP_Temp, lambda1, lambda3 and D-TAT show them to be the major factors, and agree with our regression tree results. Note that the numbers of spares are not as important as other factors. This means that raising the spare levels from their lowest levels to the highest levels in Table 1 does not lead to appreciable improvement in the average operational availability.

Sorted Parameter Estimates

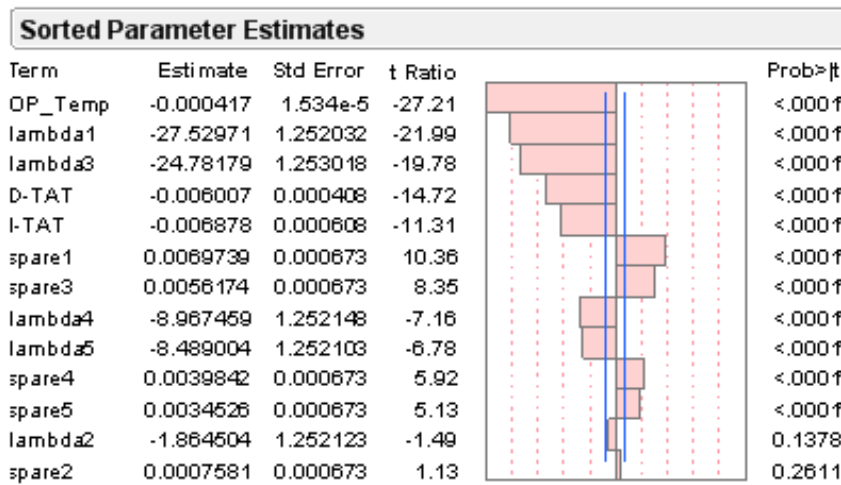| Term	Estimate	Std Error	t Ratio		Prob>|t|
OP_Temp	-0.000417	1.534e-5	-27.21		<.0001
lambda1	-27.52971	1.252032	-21.99		<.0001
lambda3	-24.78179	1.253018	-19.78		<.0001
D-TAT	-0.006007	0.000408	-14.72		<.0001
I-TAT	-0.006878	0.000608	-11.31		<.0001
spare1	0.0069739	0.000673	10.36		<.0001
spare3	0.0056174	0.000673	8.35		<.0001
lambda4	-8.967459	1.252148	-7.16		<.0001
lambda5	-8.489004	1.252103	-6.78		<.0001
spare4	0.0039842	0.000673	5.92		<.0001
spare5	0.0034526	0.000673	5.13		<.0001
lambda2	-1.864504	1.252123	-1.49		0.1378
spare2	0.0007581	0.000673	1.13		0.2611

Figure 5. Sorted Parameter Estimates for Ao Model (main effects only)

We use the same methodology to analyze the results from the life cycle cost model. We first observe the regression tree as shown in Figure 6. Once again the dominant factor is the operational tempo (OP_Temp). For example, the first split indicates that the average life cycle cost is $63.5 million across the 114 scenarios that had an operational tempo greater than or equal to 523 hours per year. In contrast, the average life cycle cost is $43 million among the 143 scenarios that had an operational tempo less than 523 hours per year. As the operational tempo increases, the operating cost obviously increases, and this factor has more impact on the life cycle cost than any others. In the second split, when operational tempo is greater than 523, the reliability of the driving differential (Part #3) becomes a critical factor.

We then develop regression metamodels with main and two-way interaction effects. After observing that the interactions are not significant compared to the main effects, we build a simple linear regression model, and the results are shown in Figure 7. The R-square value for this model was 0.98, demonstrating that this linear regression model does an excellent job of explaining the variability in the life cycle cost model. Again we see that operational tempo is the most critical factor followed by lambda3 and lambda1 (the failure rates of Part #3 and #1). However, the repair turnaround times (I-TAT and D-TAT) do not have much effect on the life cycle cost. As the repair turnaround times increase, the operational availability deteriorates, yet the operations and maintenance cost will not be affected. However, there is an implied waste of money since the repair pipeline inventory will go up as the repair turnaround time increases.

All Rows
Count	257	**LogWorth**	**Difference**
Mean	52125320	67.285105	2.05e+7
Std Dev	14641807		

OP_Temp<523
Count	143	**LogWorth**	**Difference**
Mean	43036017	20.757715	1.15e+7
Std Dev	9527252.3		

OP_Temp>=523
Count	114	**LogWorth**	**Difference**
Mean	63526814	27.663572	1.64e+7
Std Dev	11658664		

lambda3<0.00228
Count	61	**LogWorth**	**Difference**
Mean	36428476	12.247506	9376652
Std Dev	6788859.3		

lambda3>=0.00228
Count	82	**LogWorth**	**Difference**
Mean	47951382	24.08513	1.26e+7
Std Dev	8224769		

lambda3<0.00205
Count	46
Mean	53745169
Std Dev	6728997.6

lambda3>=0.00205
Count	68
Mean	70143809
Std Dev	9418327.8

OP_Temp<367
Count	23
Mean	30587283
Std Dev	4453484.4

OP_Temp>=367
Count	38
Mean	39963935
Std Dev	5375745.9

OP_Temp<445
Count	54
Mean	43656595
Std Dev	5476631.5

OP_Temp>=445
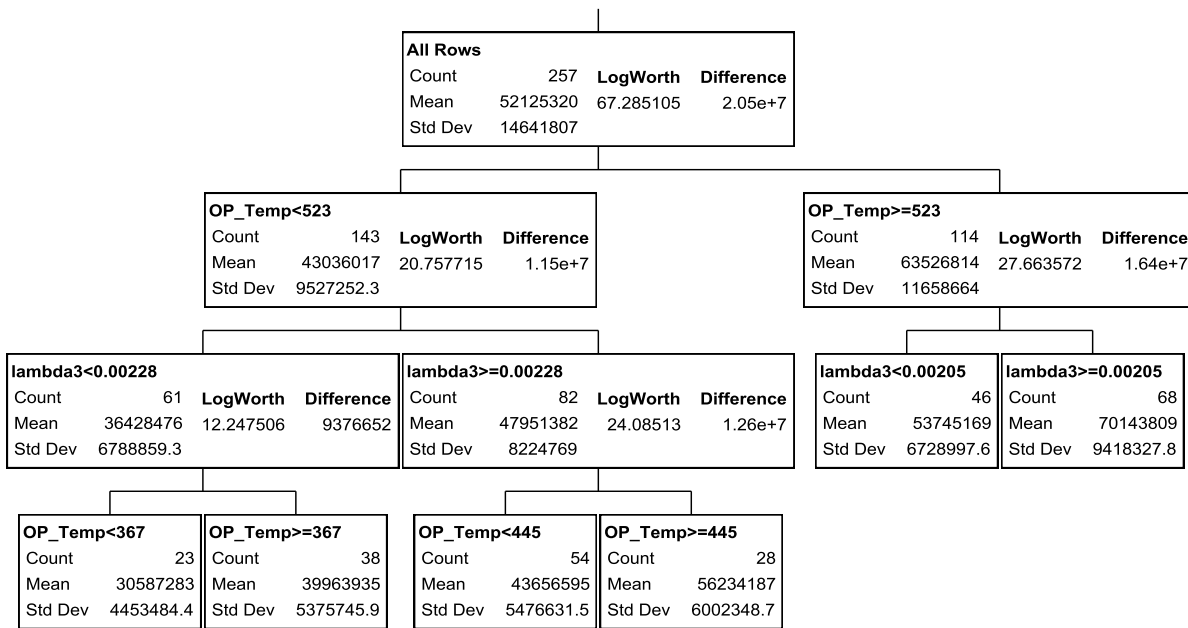| Count	28
Mean	56234187
Std Dev	6002348.7

Figure 6. Regression Tree for the Life Cycle Cost

The number of spares for Part #1 (the most expensive one among the five parts selected for analysis) seems to be more critical than other spares, yet the number of spares in general does not have much impact on the life cycle cost.

Sorted Parameter Estimates

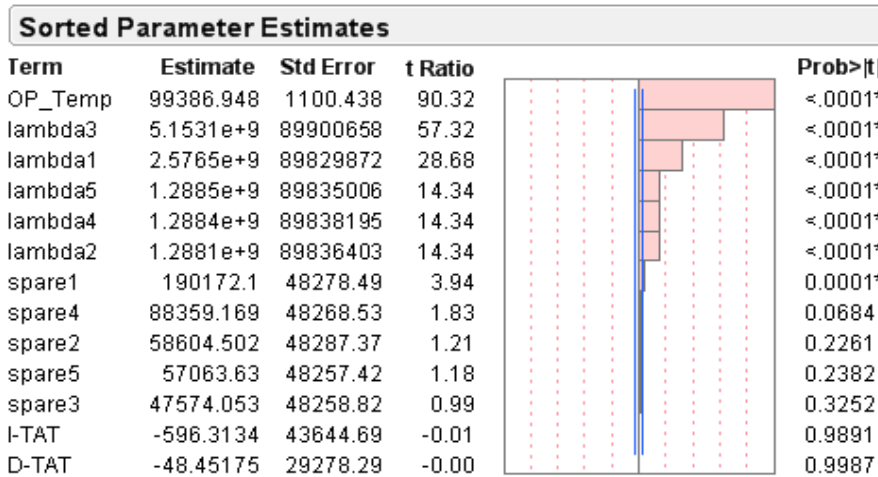| Term	Estimate	Std Error	t Ratio		Prob>|t|
OP_Temp	99386.948	1100.438	90.32		<.0001*
lambda3	5.1531e+9	89900658	57.32		<.0001*
lambda1	2.5765e+9	89829872	28.68		<.0001*
lambda5	1.2885e+9	89835006	14.34		<.0001*
lambda4	1.2884e+9	89838195	14.34		<.0001*
lambda2	1.2881e+9	89836403	14.34		<.0001*
spare1	190172.1	48278.49	3.94		0.0001*
spare4	88359.169	48268.53	1.83		0.0684
spare2	58604.502	48287.37	1.21		0.2261
spare5	57063.63	48257.42	1.18		0.2382
spare3	47574.053	48258.82	0.99		0.3252
I-TAT	-596.3134	43644.69	-0.01		0.9891
D-TAT	-48.45175	29278.29	-0.00		0.9987

Figure 7. Sorted Parameter Estimates for the Life Cycle Cost Model (main effects only)

Lastly, we compare Ao and life cycle cost together in a scatterplot, given in Figure 8. The darkened points correspond to those scenarios with lowest cost and highest Ao. There is an obvious negative correlation between the two measures, with higher operational availability generally leading to lower overall life cycle cost. In examining the input factor distributions for these selected points (not shown here), we find that the factor with the most explanatory power for achieving both lower cost and higher Ao together is operational tempo. The next most driving influence for both measures considered together is the failure rate of Part # 3. These darkened points are also shown as the darkened areas in the distribution graphs of Figures 2 and 3.
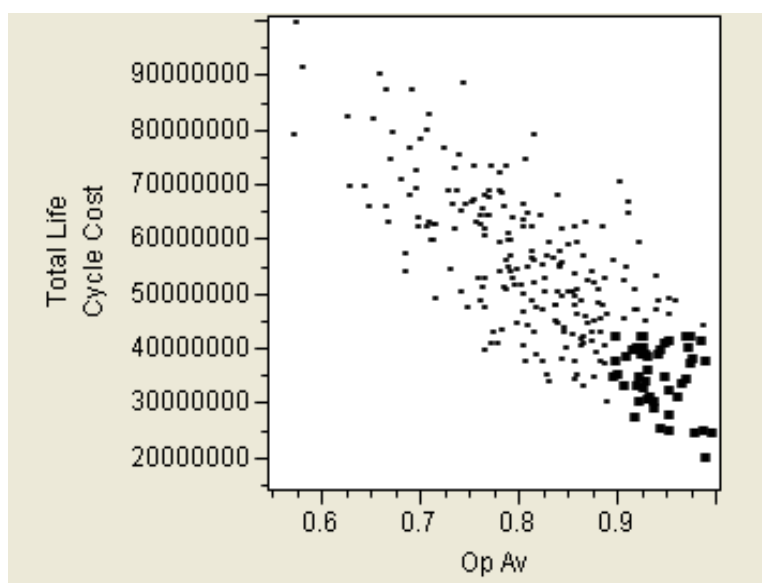
Figure 8. Scatterplot of Cost vs Ao

5   REMARKS

The most critical factor is the operational tempo for both Ao and the life cycle cost. The more you operate, the more it fails and the more it costs. Then the reliability plays an important role. Those parts with high failure rates are more critical (e.g., Part #1 and Part #3). D-level turnaround times as well as I-level turnaround times are also critical factors for Ao. Those with longer repair turnaround times [e.g., Part #1 fails less often than Part #3, yet its repair turnaround time is much longer (D-level maintenance) than that of Part #3 (I-level maintenance)] tend to be more critical for Ao. The repair turnaround times do not have much direct impact on the life cycle cost, since, in the life cycle spreadsheet model, the spare level for each component was predetermined for each scenario. However, in reality if the repair turnaround time gets longer, managers tend to purchase more spares to improve readiness, yet those newly acquired spares will only spike the operational availability for a short term. Eventually those newly acquired spares will fail and get stuck in the repair pipeline without improving Ao in the long run if the repair turnaround time is not reduced.

In conclusion, both warfighters and financial managers should understand the importance of logistics and impact of reliability, cycle time and operational tempo on readiness and life cycle cost. Our approach is intended to serve as a basis for discussion between program offices concerned with cost and operational commands concerned with operational availability.

REFERENCES

Cioppa, T.M. and T.W. Lucas. 2007. Efficient nearly orthogonal and space-filling Latin hypercubes. *Technometrics* 49(1):45-55.

Kang, K., K.H. Doerr, U. Apte, and M. Boudreau. 2010. Decision support for valuing improvements in component reliability and maintenance. To appear in *Military Operations Research*.

Kang, K., K.H. Doerr, and S.M. Sanchez. 2006. A design of experiments approach for readiness risk analysis. In *Proceedings of the 2006 Winter Simulation Conference,* ed. L.F. Perrone, F.P. Wieland, J. Liu, B.G. Lawson, D.M. Nicol, and R.M. Fujimoto, 1332-1339. Piscataway, NJ: Institute of Electrical and Electronic Engineers, Inc.

Kleijnen, J.P.C., S.M. Sanchez, T.W. Lucas, and T.M. Cioppa. 2005. A user's guide to the brave new world of designing simulation experiments. *INFORMS Journal on Computing* 17 (3): 263-289 (with online companion).

Kelton, W.D., R.P. Sadowski, and N.B. Swets. 2010. *Simulation with Arena*. Fifth Edition, McGraw-Hill.

NPS 2010. SEED Center for Data Farming. Naval Postgraduate School. Available via: <http://harvest.nps.edu> [accessed April 14, 2010].

Sanchez, S. M. 2006. Work smarter, not harder: guidelines for designing simulation experiments. *Proceedings of the 2006 Winter Simulation Conference,* eds. L.F. Perrone, F.P. Wieland, J. Liu, B.G. Lawson, D.M. Nicol, and R.M. Fujimoto, 47-57. Piscataway, NJ: Institute of Electrical and Electronic Engineers, Inc.

SAS. 2008. *JMP® User's Guide, Version 8*. Cary, NC: SAS Institute Inc.

AUTHOR BIOGRAPHIES

KEEBOM KANG is on the faculty of the Graduate School of Business and Public Policy of the Naval Postgraduate School (NPS) in Monterey, California. He served as the Chair of the Operations and Logistics program for 5 years (2004-2009). His research interests are in the areas of simulation and military logistics analysis. He received his Ph.D. in Industrial Engineering from Purdue University; an M.S. in Operations Research from the University of Texas at Austin; and a B.S. in Industrial Engineering from Seoul National University. He is currently Associate Editor of *Naval Research Logistics*. His email address is <kkang@nps.edu>.

MARY MCDONALD is a Research Associate at the SEED Center for Data Farming at the Naval Postgraduate School. She has a B.A in Mathematics from Northwestern University and M.S. in Applied Mathematics from the Naval Postgraduate School. She has 15 years of experience in military simulation and analysis in both academia and industry. She previously served 11 years on active duty with the U.S. Marine Corps in the fields of Financial Management and Applied Mathematics, and is currently a Lieutenant Colonel in the Reserves. Her email address is <mlmcdona@nps.edu>.